

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Childress et al.	§	
	§	Group Art Unit: 2444
Serial No.: 10/753,858	§	
	§	Examiner: Anwari, Maceeh
Filed: January 8, 2004	§	
	§	Confirmation No.: 6042
For: Method and Apparatus for Non-	§	
Invasive Discovery of Relationships	§	
Between Nodes in a Network		

35525

PATENT TRADEMARK OFFICE
CUSTOMER NUMBER

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF (37 C.F.R. 41.37)

This brief is in furtherance of the Notice of Appeal, filed in this case on January 8, 2009.

A fee of \$540.00 is required for filing an Appeal Brief. Please charge this fee to IBM Corporation Deposit Account No. 09-0447. No additional fees are believed to be necessary. If, however, any additional fees are required, I authorize the Commissioner to charge these fees which may be required to IBM Corporation Deposit Account No. 09-0447.

REAL PARTY IN INTEREST

The real party in interest in this appeal is the following party: International Business Machines Corporation of Armonk, New York.

RELATED APPEALS AND INTERFERENCES

This appeal has no related proceedings or interferences.

STATUS OF CLAIMS

A. TOTAL NUMBER OF CLAIMS IN APPLICATION

The claims in the application are: 1-21

B. STATUS OF ALL THE CLAIMS IN APPLICATION

Claims canceled: 1-8

Claims withdrawn from consideration but not canceled: None

Claims pending: 9-21

Claims allowed: None

Claims rejected: 9-21

Claims objected to: None

C. CLAIMS ON APPEAL

The claims on appeal are: 9-21

STATUS OF AMENDMENTS

No amendment after final rejection was filed.

SUMMARY OF CLAIMED SUBJECT MATTER

A. CLAIM 9 - INDEPENDENT

The subject matter of claim 9 is directed to a data processing system for identifying nodes in a network data processing system as described at least in the specification on page 4, lines 3-5; page 7, line 3 to page 11, line 19; Figure 1, item 100; Figure 2, item 200; and Figure 3, item 300. The data processing system comprises a bus system, a communications unit connected to the bus system, a memory connected to the bus system, wherein the memory includes a set of instructions, and a processing unit connected to the bus system as described at least in the specification on page 8, line 9 to page 11, line 19, Figure 2, item 200, and Figure 3, item 300. The processing unit as described at least in the specification on page 9, line 30 to page 10, line 1; page 10, lines 19-21; and Figure 3, item 302 executes the set of instructions to receive cache data from a set of routers in the data processing system on a periodic basis, wherein the cache data includes an identification of the nodes sending data packets onto the network data processing system as described at least on page 12, line 23 to page 13, line 2; Figure 4, items 400, 408, 410, 412, 414; and page 13, lines 17-25. Responsive to receiving the cache data, the cache data is stored prior to clearing the cache data present in the set of routers, wherein the stored cache data comprises snapshots of cache data previously present in the set of routers over time as described at least on page 13, lines 9-16; and Figure 4, items 400, 404, 412. The processing unit identifies the nodes on the network data processing system using the stored cache data from the set of routers as described at least on page 13, lines 9-16; Figure 4, blocks 406, 416, 418; page 17, lines 6-28; and Figure 8, items 800-808. The processing unit then generates a display of the nodes in a graphical view comprising communications paths between the nodes with a graphical indication of network traffic volume using the stored cache data, wherein the graphical view includes network traffic volume and node relationships over time as described at least on page 13, line 26 to page 14, line 2; page 15, lines 3-28; and Figure 6, item 600.

B. CLAIM 10 - INDEPENDENT

The subject matter of claim 12 is directed to a data processing system for identifying nodes in a network data processing system as described at least in the specification on page 4, lines 3-5; page 7, line 3 to page 11, line 19; Figure 1, item 100; Figure 2, item 200; and Figure 3,

item 300. The data processing system comprises a receiving means for receiving cache data from a set of routers in the data processing system on a periodic basis, wherein the cache data includes an identification of the nodes sending data packets onto the network data processing system as described at least on page 12, line 23 to page 13, line 2; Figure 3, item 302; Figure 4, items 400, 408, 410, 412, 414; and page 13, lines 17-25. The data processing system further comprises storing means for, in response to receiving the cache data, storing the cache data prior to clearing the cache data present in the set of routers, wherein the stored cache data comprises snapshots of cache data previously present in the set of routers over time as described at least on page 13, lines 9-16; Figure 3, item 302; and Figure 4, items 400, 412, 404, 416. The data processing system further comprises identifying means for identifying the nodes on the network data processing system using the stored cache data from the set of routers as described at least on page 13, lines 9-16; Figure 3, item 302; Figure 4, blocks 406, 416, 418; page 17, lines 6-28; and Figure 8, items 800-808. The data processing system further comprises generating means for generating a display of the nodes in a graphical view comprising communications paths between the nodes with a graphical indication of network traffic volume using the stored cache data, wherein the graphical view includes network traffic volume and node relationships over time as described at least on page 13, line 26 to page 14, line 2; page 15, lines 3-28; Figure 3, item 302; Figure 4, blocks 406, 416, 418; and Figure 6, item 600.

C. CLAIM 12 - DEPENDENT

The subject matter of claim 12 is directed to the data processing system of claim 10 further comprising identifying means for identifying communications paths between the nodes on the network data processing system using the stored cache data as described at least on page 13, line 30 to page 14, line 2; Figure 3, item 302; page 15, lines 3-28; and Figure 6, item 600.

D. CLAIM 18 - INDEPENDENT

The subject matter of claim 12 is directed to a computer program product in a computer readable medium for identifying nodes in a network data processing system as described at least in the specification on page 18, line 19 to page 19, line 7; and page 4, lines 3-5. The computer program product comprises first instructions for receiving cache data from a set of routers in the data processing system on a periodic basis, wherein the cache data includes an identification of

the nodes sending data packets onto the network data processing system as described at least on page 12, line 23 to page 13, line 2; Figure 4, items 400, 408, 410, 412, 414; and page 13, lines 17-25. The computer program product further comprises second instructions for, in response to receiving the cache data, storing the cache data prior to clearing the cache data present in the set of routers, wherein the stored cache data comprises snapshots of cache data previously present in the set of routers over time as described at least on page 13, lines 9-16; and Figure 4, items 400, 412, 404. The computer program product further comprises third instructions for identifying the nodes on the network data processing system using the stored cache data from the set of routers as described at least on page 13, lines 9-16; Figure 4, blocks 406, 416, 418; page 17, lines 6-28; and Figure 8, items 800-808. The computer program product further comprises fourth instructions for generating a display of the nodes in a graphical view comprising communications paths between the nodes with a graphical indication of network traffic volume using the stored cache data, wherein the graphical view includes network traffic volume and node relationships over time as described at least on page 13, line 26 to page 14, line 2; page 15, lines 3-28; and Figure 6, item 600.

GROUND OF REJECTION TO BE REVIEWED ON APPEAL

The ground of rejection to review on appeal is as follows:

A. GROUND OF REJECTION 1

Whether claims 9-21 were properly rejected under 35 U.S.C. § 103 as being unpatentable over Lam (U.S. Publication No. 2007/0113017) and further in view of Pruthi et al. (U.S. Publication No. 2002/0105911).

ARGUMENT

A. GROUND OF REJECTION 1 (Claims 9-21)

The Examiner rejected claims 9-21 under 35 U.S.C. § 103 as being unpatentable over Lam (U.S. Publication No. 2007/0113017), hereinafter “*Lam*” and further in view of Pruthi et al. (U.S. Publication No. 2002/0105911), hereinafter “*Pruthi*”.

The Examiner states:

Regarding claim 9, Lam discloses: A data processing system for identifying nodes in a network data processing system, the data processing system comprising: a bus system (Figure 7 and par. 56, 58; SCSI bus); a communications unit connected to the bus system (Figures 1-2, 7-8 and par. 27; computer, processor, various servers and network); a memory connected to the bus system, wherein the memory includes a set of instructions (Figures 1-2, 7-8 and par. 27; [210, 250-1 324,340,245,325,719,725,732] generating commands); and a processing unit connected to the bus system, wherein the processing unit executes the set of instructions to receive cache data from a set of routers in the data processing system on a periodic basis, wherein the cache data includes an identification of the nodes sending data packets onto the network data processing system (Figures 1-2, 5, 7-9 and Abstract & par. 27 &37; [508-540]recording data, computer/processor and cache data); in response to receiving the cache data, store the cache data basis prior to clearing the cache data present in the set of routers, wherein the stored cache data comprises snapshots of cache data previously present in the set of routers over time (Figures 5 &9 and Abstract; [508-540] [719-750] determining requirement of snapshot); identify the nodes on the network data processing system using the stored cache data from the set of routers (par. 34, 37-38, 45 and 56; identifiers); and generate a display of the nodes in a graphical view comprising communications paths between the nodes with a graphical indication of network traffic volume using the stored cache data, wherein the graphical view includes network traffic volume and node relationships over time.

Lam does not explicitly teach generating a display of the nodes in a graphical view comprising communications paths between the nodes with a graphical indication of network traffic volume using the stored cache data, wherein the graphical view includes network traffic volume and node relationships over time.

In the same field of endeavor, Pruthi teaches generating a display of the nodes in a graphical view comprising communications paths between the nodes with a graphical indication of network traffic volume using the stored cache data, wherein the graphical view includes network traffic volume and node relationships over time (Figures 10- 23 and Abstract; traffic plots).

Accordingly it would have been obvious for one of ordinary skill in the art to modify or incorporate Pruthi's teachings of graphically depicting network traffic

and node relationships over time with the teachings of Lam, to collect, analyze and monitor data over a communications network (Pruthi par. 2).

Final Office Action dated October 8, 2008, pages 2-4.

The Examiner bears the burden of establishing a *prima facie* case of obviousness based on the prior art when rejecting claims under 35 U.S.C. § 103. *In re Fritch*, 972 F.2d 1260, 23 U.S.P.Q.2d 1780 (Fed. Cir. 1992). For an invention to be *prima facie* obvious, the prior art must teach or suggest all claim limitations. *In re Royka*, 490 F.2d 981, 180 USPQ 580 (CCPA 1974). Independent claim 9, which is representative of independent claims 10 and 18 with regard to similarly recited subject matter, reads as follows:

9. A data processing system for identifying nodes in a network data processing system, the data processing system comprising:
a bus system;
a communications unit connected to the bus system;
a memory connected to the bus system, wherein the memory includes a set of instructions; and
a processing unit connected to the bus system, wherein the processing unit executes the set of instructions to receive cache data from a set of routers in the data processing system on a periodic basis, wherein the cache data includes an identification of the nodes sending data packets onto the network data processing system; in response to receiving the cache data, store the cache data prior to clearing the cache data present in the set of routers, wherein the stored cache data comprises snapshots of cache data previously present in the set of routers over time; identify the nodes on the network data processing system using the stored cache data from the set of routers; and generate a display of the nodes in a graphical view comprising communications paths between the nodes with a graphical indication of network traffic volume using the stored cache data, wherein the graphical view includes network traffic volume and node relationships over time.

Claim 9 recites the features of receiving cache data from the routers on a periodic basis. Claim 9 also recites that when cache data is received (on the periodic basis), this cache data is stored. By receiving cache data periodically, storing this periodically-received cache data, and then clearing the cache(s) each time cache data is received, the stored cache data comprises snapshots of the routers' cache data at various points in time. Additionally, the claim recites that the caches at the routers are cleared when cache data is received from the routers.

Lam and *Pruthi*, either alone or in combination, do not teach or suggest receiving cache

data from a set of routers in the data processing system on a periodic basis, wherein the cache data includes an identification of the nodes sending data packets onto the network data processing system as recited in amended claim 9 of the present invention. The Examiner relies on the *Lam* reference as teaching this feature. *Lam* is directed to a system for performing snapshot copies of data stored in a data storage system in a manner that captures data transmitted to disk by a client server prior to the snapshot request without requiring the client server to cease sending data to the storage system. (*Lam*, paragraph [0010].) The sections of *Lam* cited by the Examiner disclose a general network environment components of a storage system, and SCSI bus phases in which *Lam* may be implemented, as well as the process for performing a snapshot of data in a storage device. For example, cited Figure 5, blocks 508-540 disclose that a controller determines whether a snapshot should be made based on either receiving a snapshot request from a client server or due to normal operation of the system, and if so, the controller generates and inserts a snapshot marker that is used to identify the last item of data that is to be sent to disk before the snapshot is performed into the cache at the end of the data queue. The controller then flushes the data from the cache to disk as long as the snapshot marker is in the cache. Once the snapshot marker is flushed from the cache, the controller performs the requested snapshot of the data. In the Abstract, *Lam* discloses that the cache receives initial data items, these initial data items are flushed to a storage device, and the initial data items are recorded. As the initial data items are being flushed to the storage device, additional data items may be received by the cache. In paragraph [0038], *Lam* discloses that disk space on a virtual storage device is divided into segments having associated segment descriptors. These storage descriptors contain information defining the segment it represents, such as a home storage location, physical starting sector of the segment, sector count within the segment, sector count within the segment, and segment number. A segment map is used to provide the logical sector to physical sector mapping of the virtual storage device.

While the cited sections of *Lam* disclose a system for recording snapshots of data flushed from a cache, *Lam* does not teach or suggest that the cache data received from a set of routers in the data processing system on a periodic basis includes an identification of the nodes sending data packets onto the network data processing system as recited in the presently claimed invention. Rather, *Lam* merely discloses that data items are received by the cache, and this data may include initial data items and additional data items. The initial data items are transmitted to

the storage system prior to a moment in time and the additional data items are transmitted to the storage system after the moment in time. *Lam* teaches that data may be received at a cache, and the initial data items may be flushed to a storage device while the cache receives additional data items. In contrast, the presently claimed invention recites when data is received from a router, the data comprises information that specifies an ID of the node that sent the data. *Lam* does not disclose that the data comprises an identification of nodes from which the data was sent, or that the data is received at the cache from a set of routers. Rather, *Lam* merely discusses that data received by the cache may be defined by when the data was received (e.g., initial data or additional data).

Lam and *Pruthi* also do not teach or suggest identifying the nodes on the network data processing system using the stored cache data from the set of routers as recited in claim 9 of the present invention. The Examiner again relies on the *Lam* reference as teaching this feature. The sections of *Lam* cited by the Examiner are reproduced below:

Each of storage devices **250** may be any type of storage device that allows block-level storage access. In one implementation, storage devices **250** are disk drives. A disk drive includes one or more disks having a medium upon which information may be written. Each disk includes a number of physical tracks, each of which, in turn, is divided into one or more physical blocks. Accordingly, in this implementation, an address identifying the location of a data block on a disk drive may specify a disk, a physical track and a physical block. Storage devices **250-1**, **250-2** and **250-3** are connected to storage system manager **220**, in accordance with this implementation, by Fibre Channel interfaces, SCSI connections, or a combination thereof.

Lam, paragraph [0034].

According to this implementation, when storage system manager **220** initially defines a virtual storage device, or when additional storage is assigned to the virtual storage device, the disk space on the storage devices is divided into segments. Each segment has associated with it segment descriptors, which are stored in a free list table in memory. Generally, a segment descriptor contains information defining the segment it represents; for example, the segment descriptor may define a home storage device location, physical starting sector of the segment, sector count within the segment, and segment number. FIG. **3A** illustrates schematically the contents of a segment descriptor **32**. Fields **32-1** through **32-5** contain data indicating, respectively, on which storage device the segment is located, the segment's starting physical address, the segment's size, a

segment identifier, and other information which may be useful for identifying and processing the segment.

Referring to FIG. **3B**, as segments are needed to store data, the next available segment descriptor, e.g., segment descriptor **32**, is identified from the free segment list **50**, the data is stored in the segment, and the segment descriptor **32** is assigned to a new table called a segment map **66**. The segment map **66** maintains information representing how each segment defines the virtual storage device. More specifically, the segment map provides the logical sector to physical sector mapping of a virtual storage device. After the free segment descriptor **32** is moved or stored in the appropriate area of the segment map **66**, which in this example is slot **2 (70)**, the descriptor is no longer a free segment but is now an allocated segment. A detailed description of this method for dynamically allocating disk space can be found in U.S. Pat. No. 7,058,788, dated Jun. 6, 2006, which is incorporated herein by reference in its entirety.

Lam, paragraph [0037]-[0038].

At step **510**, controller **320** generates a snapshot marker. The snapshot marker is information that identifies the last item of data that is to be sent to disk before the snapshot is performed. According to one implementation, the snapshot marker comprises a data pattern that is statistically unlikely to occur in transmissions that are received by storage system **210**. For example, a snapshot marker may be implemented by the data string:

“\$\$\$\$\$SNAPSHOTMARKER****.”

Lam, paragraph [0045]-[0046].

The central item of hardware in a SCSI system is the SCSI bus, which must conform to the specification of the SCSI standard. A SCSI device refers to any device that is connected to the SCSI bus. Each SCSI device on the bus is assigned a SCSI ID that uniquely identifies the device during SCSI transactions.

Lam, paragraph [0056].

Paragraph [0034] of *Lam* discloses storage devices onto which the flushed cache data may be stored. This cited paragraph discloses that storage device may be disk drives, wherein each disk includes a number of physical tracks that are divided into one or more physical blocks. An address identifying the location of a data block on a disk drive may specify a disk, a physical track, and a physical block.

As stated above, paragraphs [0037]-[0038] of *Lam* disclose that disk space on a virtual

storage device is divided into segments having associated segment descriptors. These storage descriptors contain information defining the segment it represents, such as a home storage location, physical starting sector of the segment, sector count within the segment, sector count within the segment, and segment number. A segment map is used to provide the logical sector to physical sector mapping of the virtual storage device.

Paragraphs [0045]-[0046] of *Lam* disclose a snapshot marker that is used to identify the last item of data that is to be sent to disk before the snapshot is performed into the cache at the end of the data queue. While *Lam* mentions that the snapshot marker identifies the last item of data to be sent to disk before the snapshot is taken, *Lam* does not disclose an identification of the particular node that sent the data item.

Paragraph [0056] of *Lam* teaches SCSI devices on the SCSI bus is assigned a SCSI ID that uniquely identifies that device during SCSI transactions. As disclosed in *Lam*, the SCSI ID is utilized during an arbitration phase of a SCSI transaction in that SCSI devices assert their IDs by transmitting a signal onto the bus. The SCSI device having the highest ID wins the arbitration and becomes the initiator for the next SCSI transaction. However, there is no mention in *Lam* of an identifying process that uses the cache data now stored in a storage device to identify the nodes on the network data processing system, or that the nodes identified in the identifying process are used to generate a display of the nodes in a graphical view. Instead, the cited portions of *Lam* only teach an address identifying the location of a data block on a disk and the use of SCSI IDs for arbitration purposes, rather than the identification, from the stored cache data, of particular nodes that sent the data packets.

Additionally, the Examiner failed to state a *prima facie* obviousness rejection against claim 1 because the Examiner failed to state a proper reason to combine the references under the standards of *KSR Int'l*. Rejections on obviousness grounds cannot be sustained by mere conclusory statements; instead, there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness. *KSR Int'l. Co. v. Teleflex, Inc.*, No. 04-1350 (U.S. Apr. 30, 2007). (citing *In re Kahn*, 441 F.3d 977, 988 (CA Fed. 2006)).

Regarding a reason to combine the references, the Examiner states:

...it would have been obvious to one of ordinary skill in the art to modify or incorporate Pruthi's teachings of graphically depicting network traffic and node relationships over time with the teachings of *Lam* to collect, analyze and monitor

data over a communications network.

Final Office Action dated October 8, 2008, page 4.

However, this reason is not a rational underpinning to support the legal conclusion of obviousness of the claim in view of the combination of the references when considered as a whole. The Examiner fails to state why one of ordinary skill in the art would look to modify *Pruthi*'s method of providing a network monitoring system for collecting and analyzing communication data (*Pruthi*, paragraphs [0002] and [0016] with *Lam*'s method of performing snapshot copies of data stored in a storage system in a manner that captures data transmitted to disk by a client server prior to the snapshot request without requiring the client server to cease sending data to the storage system (*Lam*, paragraph [0010]). *KSR Int'l.* requires that the Examiner provide a rational underpinning to support the legal conclusion of obviousness. By failing to state why one of ordinary skill in the art would look to modify the method as taught by the *Pruthi* reference with the method as taught by the *Lam* reference, the Examiner has failed to provide a basis to support the legal conclusion of obviousness. Instead, the Examiner has only provided a conclusory statement and then assumed the legal conclusion without the required analysis. Therefore, the Examiner's statement does not provide a rational underpinning to support the legal conclusion of obviousness, as required by *KSR Int'l.*

In addition, it would not have been obvious to one of ordinary skill in the art to combine *Pruthi* and *Lam* since *Pruthi* would not require or benefit from the use of performing snapshot copies of data in a storage system in *Lam* to collect and analyze network traffic. *Pruthi* is directed to providing a network monitoring system for collecting and analyzing communication data, rather than ensuring that a client server can still send data to the storage system while a snapshot of data in the storage system is being performed. In addition, *Pruthi* would not require or benefit from the use of the snapshot data in *Lam* to collect and analyze network traffic. As there would be no reason for *Pruthi* to perform snapshot copies of data in a storage system, there would be no reason for one of ordinary skill in the art to have combined the teaching of *Pruthi* with the teachings of *Lam*. One would receive no benefit of such a combination. Consequently, one of ordinary skill would not be motivated by the teachings of *Pruthi* and *Lam* to combine the references to reach the presently claimed invention. As a result, the Examiner failed to state a *prima facie* obviousness rejection against claims 1, 10, and 18.

Consequently, as *Lam* and *Pruthi*, either alone or in combination, fail to teach or suggest all of the features of claims 9, 10, and 18, the rejection of claims 9, 10, and 18 has been overcome.

In addition, since claims 11-17 and claim 21 depend from claim 10, and claims 19-20 depend from claim 18, the same distinctions between *Lam* and *Pruthi*, and the claimed invention in claims 11-17 and 19-21 apply for these dependent claims. Furthermore, these dependent claims include additional features not found in the cited reference.

For example, claim 12 recites identifying means for identifying communications paths between the nodes on the network data processing system using the stored cache data. The Examiner alleges that this feature is taught by the ‘identifiers’ disclosed in paragraphs [0034], [0037]-[0038], [0045], and [0056] of *Lam*. However, as discussed above, the identifiers in *Lam* are disclosed as either an address that is used to identify the location of a data block on a disk, or a SCSI ID for identifying which SCSI device has a higher ID for arbitration purposes. In contrast, the presently claimed invention recites identifying, from the stored cache data, the communication paths of the nodes that sent the data. *Lam* only discloses using identifiers to identify a device for arbitration purposes or to identify space on a disk. Consequently, *Lam* cannot teach or suggest identifying communications paths between the nodes on the network data processing system using the stored cache data.

Therefore, the rejection of claims 9-21 under 35 U.S.C. § 103 has been overcome.

B. CONCLUSION

As shown above, the Examiner has failed to state valid rejections against any of the claims. Therefore, Appellants request that the Board of Patent Appeals and Interferences reverse the rejections. Additionally, Appellants request that the Board direct the Examiner to allow the claims.

Date: February 27, 2009

Respectfully submitted,

/Cathrine K. Kinslow/

Cathrine K. Kinslow
Reg. No. 51,886
Yee & Associates, P.C.
P.O. Box 802333
Dallas, TX 75380
(972) 385-8777

CLAIMS APPENDIX

The text of the claims involved in the appeal is as follows:

9. A data processing system for identifying nodes in a network data processing system, the data processing system comprising:

a bus system;

a communications unit connected to the bus system;

a memory connected to the bus system, wherein the memory includes a set of instructions; and

a processing unit connected to the bus system, wherein the processing unit executes the set of instructions to receive cache data from a set of routers in the data processing system on a periodic basis, wherein the cache data includes an identification of the nodes sending data packets onto the network data processing system; in response to receiving the cache data, store the cache data prior to clearing the cache data present in the set of routers, wherein the stored cache data comprises snapshots of cache data previously present in the set of routers over time; identify the nodes on the network data processing system using the stored cache data from the set of routers; and generate a display of the nodes in a graphical view comprising communications paths between the nodes with a graphical indication of network traffic volume using the stored cache data, wherein the graphical view includes network traffic volume and node relationships over time.

10. A data processing system for identifying nodes in a network data processing system, the data processing system comprising:

receiving means for receiving cache data from a set of routers in the data processing system on a periodic basis, wherein the cache data includes an identification of the nodes sending data packets onto the network data processing system;

storing means for, in response to receiving the cache data, storing the cache data prior to clearing the cache data present in the set of routers, wherein the stored cache data comprises snapshots of cache data previously present in the set of routers over time;

identifying means for identifying the nodes on the network data processing system using the stored cache data from the set of routers; and

generating means for generating a display of the nodes in a graphical view comprising communications paths between the nodes with a graphical indication of network traffic volume using the stored cache data, wherein the graphical view includes network traffic volume and node relationships over time.

11. The data processing system of claim 10, wherein the cache data is from a set of address resolution protocol caches located on the set of routers.

12. The data processing system of claim 10 further comprising:

identifying means for identifying communications paths between the nodes on the network data processing system using the stored cache data.

13. The data processing system of claim 11, further comprising:
identifying means for identifying network traffic on the communication paths using the stored cache data.
14. The data processing system of claim 13 , wherein the stored cache data is used to validate service level agreement compliance.
15. The data processing system of claim 11, wherein the cache data is received through agents located on the set of routers.
16. The data processing system of claim 15, where the agents clear the set of address resolution protocol caches each time data is sent to the data processing system.
17. The data processing system of claim 11, wherein the cache data contains entries for the nodes sending data packets onto the network data processing system and wherein each entry includes at least one of a media access control address, a source Internet Protocol address, and a destination Internet Protocol address.
18. A computer program product in a computer readable medium for identifying nodes in a network data processing system, the computer program product comprising:
first instructions for receiving cache data from a set of routers in the data processing system on a periodic basis, wherein the cache data includes an identification of the nodes sending data packets onto the network data processing system;

second instructions for, in response to receiving the cache data, storing the cache data prior to clearing the cache data present in the set of routers, wherein the stored cache data comprises snapshots of cache data previously present in the set of routers over time;

third instructions for identifying the nodes on the network data processing system using the stored cache data from the set of routers; and

fourth instructions for generating a display of the nodes in a graphical view comprising communications paths between the nodes with a graphical indication of network traffic volume using the stored cache data, wherein the graphical view includes network traffic volume and node relationships over time.

19. The computer program product of claim 18, wherein the cache data is from a set of address resolution protocol caches located on the set of routers.

20. The computer program product of claim 18 further comprising:

fifth instructions for identifying communications paths between the nodes on the network data processing system using the stored cache data.

21. The data processing system of claim 10, wherein the graphical indication comprises network connections of different thicknesses to indicate network traffic volume.

EVIDENCE APPENDIX

This appeal brief presents no additional evidence.

RELATED PROCEEDINGS APPENDIX

This appeal has no related proceedings.